

Distributed team formation in multi-agent systems: stability and approximation

Lorenzo Coviello and Massimo Franceschetti

Abstract

We consider a scenario in which leaders are required to recruit teams of followers. Each leader cannot recruit all followers, but interaction is constrained according to a bipartite network. The objective for each leader is to reach a state of *local stability* in which it controls a team whose size is equal to a given constraint. We focus on distributed strategies, in which agents have only local information of the network topology and propose a distributed algorithm in which leaders and followers act according to simple local rules. The performance of the algorithm is analyzed with respect to the convergence to a *stable solution*.

Our results are as follows. For any network, the proposed algorithm is shown to converge to an *approximate* stable solution in polynomial time, namely the leaders quickly form teams in which the total number of additional followers required to satisfy all team size constraints is an arbitrarily small fraction of the entire population. In contrast, for general graphs there can be an exponential time gap between convergence to an approximate solution and to a stable solution.

I. INTRODUCTION

A multi-agent system (MAS) is composed of many interacting intelligent agents. Agents can be software, robots, or humans, and the system is highly distributed, as agents do not have a global view of the state and act autonomously of each other. These systems can be used to collectively solve problems that are difficult to solve by a single entity. Their application ranges from robotics, to disaster response, social structures, crowd-sourcing etc. A main feature of MAS is that they can manifest self-organization as well as other complex control paradigms

This work was partially supported by the Army Research Office grant number W911NF-11-1-0363.

The authors are with the Department of Electrical and Computer Engineering, University of California San Diego, 9500 Gilman Dr., La Jolla CA, 92093. Emails: lcoviell@ucsd.edu, massimo@ece.ucsd.edu

even when the individual strategies of the agents are very simple. In short, simple local interaction can conspire to determine complex global behaviors. Examples of such emerging behaviors are in economics and game theory, where local preferences translate into global equilibria [35], in social sciences, where local exposure governs the spread of innovation [41], and in control, where local decision rules determine whether and how rapidly consensus is reached [4], [28], [29], [30], [36], [37].

From a practical perspective, the performance of a MAS often depends on how quickly convergence to a global, possibly approximate, solution is reached and it is in general influenced by the network structure. For example, in the context of information diffusion in social networks, the rate of convergence of the system’s dynamics is affected by the underlying network and the local interaction rules [21], [26].

One of the critical issues in multi-agent systems is coordination. Due to the autonomous behavior of the agents and to the absence of a central controller, coordination must be distributed. In the case of human agents, it is also important that the distributed control algorithm is simple enough to be suitable to model basic principles of human behavior [8]. Two prominent problems related to consensus and coordination in multi-agent systems are leader election and group formation. In the former case, multiple agents elect a leader that can then assign tasks [24], while in the latter they divide themselves into teams in such a way that each agent knows to what team it belongs [11]. In both cases agents are all equal and coordination occurs among agents of a single class.

We consider a scenario in which there are agents of two classes, *leaders* and *followers*. Each leader must recruit a team of followers whose size is equal to a given constraint, by sending requests to the followers. Followers can only accept or reject incoming leaders’ requests. While multiple followers can be part of a leader’s team, each follower can be part of a single team at any time, but is allowed to change team over time. Moreover, a leader cannot recruit *all* followers, but can only recruit the followers it is in direct communication with. The communication structure between leaders and followers is captured by an arbitrary bipartite network, and we assume that each agent has knowledge of and can interact with its neighbors only. That is, agents only have local knowledge of the underlying network. In general, the communication constraints of the population (and therefore the structure of the bipartite network) can be dictated by physical constraints (as for example antenna visibility range or signal to noise ratio threshold), social




Fig. 1. Example of a bipartite network between leaders and followers determined by physical constraints. Left: each leader can only recruit the followers in its visibility range (dotted circle), arrows represent team membership, and the set of arrows defines a partition of the followers into teams.

Right: the resulting bipartite network. An edge between leader ℓ and follower f exists if and only if f is in ℓ 's visibility range. Matching edges define team membership and are highlighted.

context, and so on. A pictorial representation of a bipartite network arising from physical constraints is given in Figure 1.

We consider a notion of stability in which each agent controls a team of adequate size. Each leader has an incentive to reach *local* stability (that is, to build a team of followers of the right size) by dynamically interacting with its neighbors. The question we aim to answer is: can simple local rules lead to stable, or *close* to stable, team formation in reasonable time? By “close to stable” we mean that the total number of additional followers required to satisfy all team size constraints is an arbitrary small fraction of the entire population. We propose a simple, distributed, memoryless algorithm in which leaders do not communicate between each other, and we show that, in any network of size n , any constant approximation of a stable outcome (or of a suitably defined *best* outcome if a stable one does not exist) is reached in time polynomial in n with high probability. In contrast, for general graphs we show through a counterexample that there can be an exponential gap between the time needed to reach stability and that needed to reach approximate stability, that is, to find the *best* solution compared to a *good* solution. We remark that, in its simplicity, the proposed algorithm is suitable to model human agents, it can be programmed on simple robots with limited computation abilities, and it is amenable to

analysis.

The rest of the paper is organized as follows. After discussing how our work relates to the existing literature, in Section II we formally define the problem and the notions of stability and approximate stability, in Section III we present the distributed algorithm for leaders and followers, in Sections IV and V we present our technical results on the algorithm’s performance, and in Section VI we further discuss the algorithm’s performance by showing some simulations’ results. To prove our result on the convergence to approximate stability, we derive a technical lemma (Lemma 1) that relates the quality of a matching to the existence of particular paths (that we call *deficit-decreasing* paths) of given length. The lemma extends a known combinatorial result by Hopcroft and Karp [15] to the setup of many-to-one matching, and can be considered to be of independent interest.

A. Related work

The problem of team formation that we consider is an example of distributed many-to-one matching in bipartite networks [2], [14], [34]. The one-to-one case has been previously studied in the context of theoretical computer science [23] [32]. In the control literature, our work is related to the distributed assignment problem and to group formation in MAS. In this framework, Moore and Passino [27] proposed a variant of the distributed auction algorithm for the assignment of mobile agents to tasks. Cenedese et al. [6] proposed a variant of the Stable Marriage algorithm [12] to solve the distributed task assignment problem. Abdallah and Lesser [1] proposed an “almost” distributed algorithm for coalition formation, allowing for a special agent with the role of “manager”. Gatson and den Jardins [13] studied a scenario of group formation where agents can adapt to the network structure. Tosić and Agha [39] proposed an algorithm for group formation based on the distributed computation of maximal cliques in the underlying network. Further work studied team formation in multi-robot systems [40], in the case where communication between agents is not allowed [3]. Other authors considered MAS composed by leaders and followers. To cite a few, Tanner [38] derived a necessary and sufficient condition for a group of interconnected agents to be controllable by one of them acting as a leader; Rahmani et al. [33] studied the controlled agreement problem in networks in which certain agents have leader roles, translating graph-theoretic properties into control-theoretic properties; Pasqualetti et al. [31] analyzed the problem of driving a group of mobile agents, represented by a network

of leaders and followers, in which follower act according to a simple consensus rule.

We distinguish ourselves from all mentioned papers, as we propose a fully distributed algorithm for group formation on arbitrary networks in which agents act according to simple local rules and perform very limited computation, and we derive performance guarantees in the form of theorems. For an exhaustive overview on distributed algorithms in multi-agent systems, the interested reader is referred to the books by Lynch [24] and by Bullo et al. [5] and the references therein, while the survey by Horling and Lesser [16] offers an overview on three decades of research on organizational paradigms as team and coalition formation.

A more recent line of research aims to study how humans connected over a network solve tasks in a distributed fashion [8], [10], [17], [19], [20], [25]. In the work of Kearns et al. [20], human subjects positioned at the vertices of a virtual network were shown to be able to collectively reach a coloring of the network, given only local information about their neighbors. Similar papers further investigated human coordination in the case of coloring [10], [17], [25] and consensus [17], [19], with the main goal of characterizing how performance is affected by the network's structure. Using experimental data of maximum matching games performed by human subjects in a laboratory setting, Coviello et al. [8] proposed a simple algorithmic model of human coordination that allows complexity analysis and prediction.

Finally, related to our work is also the research on social exchange networks [7], [22], that considers a networked scenario in which each edge is associated to an economic value, nodes have to come to an agreement on how to share these values, and each agent can only finalize a single mutual exchange with a single neighbor. Recently, Kanoria et al. [18] proposed a distributed algorithm that reaches approximate stability in linear time. However, we consider a different setup since we allow leaders to build teams of multiple followers.

II. PROBLEM FORMULATION

We consider a population composed of agents of two different classes: leaders and followers. Each leader is required to recruit a team of followers whose size is equal to a given constraint, by sending requests to the followers. Followers can only accept or reject leaders' requests. While multiple followers can be in a leader's team, each follower can be part of a single team at a time, but is allowed to change team over time. A leader is not allowed to recruit *all* followers, but can only recruit the followers it is in direct communication with. The communication constraints of

the population are captured by a bipartite network $G = (L \cup F, E)$ whose nodes' partition is given by the set L of leaders and the set F of followers, and where there exists an edge $(f, \ell) \in E$ between follower f and leader ℓ if and only if f and ℓ can communicate between each other (see Figure 1). Let $N_\ell = \{f \in F : (f, \ell) \in E\}$ be the neighborhood of $\ell \in L$. For each $\ell \in L$, leader ℓ is required to recruit a team of c_ℓ followers, where $c_\ell \geq 1$.

Definition 1 (Matching): A subset $M \subseteq E$ is a matching of G if for each $f \in F$ there exists at most a single $\ell \in L$ such that $(\ell, f) \in M$.

The definition of matching is consistent with the fact that multiple followers can be part of a leader's team. There is a one-to-one correspondence between matchings M of G and tuples of teams $\{T_\ell(M) : \ell \in L\}$, where $T_\ell(M)$ denotes the team of leader ℓ under the matching M . We have that $T_\ell(M) = \{f \in F : (\ell, f) \in M\} \subseteq N_\ell$ for every matching M . We consider the following notion of stability.

Definition 2 (Stable matching): Given constraints c_ℓ for each $\ell \in L$, a matching M of G is stable if and only if $|T_\ell(M)| = c_\ell$ for all $\ell \in L$.

Depending on the constraints c_ℓ , a network G might not admit a stable matching. Nonetheless, given a matching of G , we are interested in assessing its *quality*. Our main result builds on the following definitions of *deficit* of a leader and deficit of a matching.

Definition 3 (Deficit of a leader): Let ℓ be a leader with constraint $c_\ell \geq 1$, and M be a matching of G . The deficit of ℓ under the matching M is

$$d_\ell(M) = \max\{0, c_\ell - |T_\ell(M)|\}.$$

Definition 4 (Deficit of a matching): Given constraints $c_\ell \geq 1$ for each $\ell \in L$, the deficit of a matching M of G is

$$d(M) = \sum_{\ell \in L} d_\ell(M) = \sum_{\ell \in L} \max\{0, c_\ell - |T_\ell(M)|\}.$$

In words, $d_\ell(M)$ is the number of additional followers leader ℓ needs to satisfy its size constraint. Similarly, $d(M)$ sums the numbers of additional followers each leader needs to satisfy its size constraint. Given a matching M , we say that a leader ℓ is *poor* if $d_\ell(M) > 0$ (that is, $|T_\ell(M)| < c_\ell$) and *stable* if $|T_\ell(M)| = c_\ell$. In this work, we do not consider the case of $|T_\ell(M)| > c_\ell$ since we assume that each leader ℓ never recruits more than c_ℓ followers simultaneously. This can be justified by the fact that recruiting additional followers might be costly.

Observe that only poor leaders contribute to $d(M)$, and that M is stable if and only if $d(M) = 0$. Given G , two matchings of G can be compared with respect to their deficit, and the best matching of G can be defined as one minimizing the deficit.

Definition 5 (Best matching): A matching M of G is a best matching of G if $d(M) \leq d(M')$ for every matching M' of G .

Observe that a stable matching is also a best matching. Moreover, if G admits a stable matching, $d(M)$ quantifies how much M differs from a stable matching of G . In general, if M^* is a best matching of G with $d(M^*) = d^*$, then, $d(M) - d^*$ tells how much M differs from a best matching of G . Given a matching M of G , the following definition provides a measure of how well M approximates a best matching of G .

Definition 6 (Approximate best matching): Fix $\varepsilon \in [0, 1]$, and let m be the number of followers in G . Let M^* be a best matching of G . Then, a matching M is a $(1 - \varepsilon)$ -approximate best matching of G if $d(M) - d(M^*) < \varepsilon m$.

When G admits a stable matching, we are interested in the notion of approximate stable matching.

Definition 7 (Approximate stable matching): Let G admit a stable matching. Fix $\varepsilon \in [0, 1]$, and let m be the number of followers in G . Then, a matching M is a $(1 - \varepsilon)$ -approximate stable matching of G if $d(M) < \varepsilon m$.

III. THE ALGORITHM

We now present a distributed algorithm for team formation. Time is divided into rounds, and each round is composed by two stages. In the first stage, each leader acts according to the algorithm in Table 1, and in the second stage each follower acts according to the algorithm in Table 2.

First consider a leader ℓ , and let M be the matching at the beginning of a given round. If ℓ is poor (that is, $|T_\ell(M)| < c_\ell$ and $|T_\ell(M)| < |N_\ell|$ (that is, ℓ is not already matched with all followers in N_ℓ) then, with probability p (where $p \in (0, 1]$ is a fixed constant), ℓ attempts to recruit an additional follower, chosen as explained below, by sending a *matching request*. An unmatched follower in N_ℓ , if any, is chosen uniformly at random; otherwise, a follower in $N_\ell \setminus T_\ell(M)$ is chosen uniformly at random. In other words, leaders always prefer to recruit followers that are

currently unmatched over matched ones. Note that a leader tries to recruit an additional follower after checking if *local stability* holds (that is, after checking if its team size is equal to c_ℓ).

Consider now a follower f . During each round, if f has incoming requests then each request is rejected independently of the others with probability $1 - q$ (where $q \in (0, 1]$ is a fixed constant). If all incoming requests are rejected, then f does not change team (if currently matched) or it remains unmatched (if currently unmatched). Otherwise, one among the active requests is chosen uniformly at random, f joins the corresponding leader, and all the other requests are discarded. For ease of presentation, we assume that a follower is equally likely to join a team when unmatched and to change team when currently matched, but all our results hold if we consider different values of q for matched and unmatched followers (and even if we consider a different value of q for each follower, as long as each value is a constant).

Table 1 Algorithm for leader $\ell \in L$

```

if  $|T_\ell(M)| < \min\{c_\ell, |N_\ell|\}$  then
  with probability  $p$  do the following
  if  $\exists$  unmatched  $f \in N_\ell$  then
    choose an unmatched follower  $f' \in N_\ell$  u.a.r.
  else
    choose a follower  $f' \in N_\ell \setminus T_\ell(M)$  u.a.r.
  end if
  send a matching request to  $f'$ 
end if

```

The proposed algorithm enjoys several properties. It is memoryless, the actions of each agent only depend on local information, and the leaders do not communicate between each other. Also, it is *self-stabilizing*, that is, once a stable matching is reached, leaders stop recruiting followers. Moreover, it is a single-stage algorithm, that is, agents never change their behavior until stability is reached. Finally, observe that the exchanged messages can be represented by a single bit.

IV. CONVERGENCE TO APPROXIMATE STABLE MATCHINGS

In this section, we only consider networks admitting stable matchings, and we show that, given any network and any constant $\varepsilon \in (0, 1)$, a $(1 - \varepsilon)$ -approximate stable matching is reached in a

Table 2 Algorithm for follower $f \in F$

```

if  $f$  has incoming requests then
  for each leader  $\ell$  requesting  $f$  do
    with probability  $1 - q$  reject  $\ell$ 's request
  end for
  if there are active requests then
    select one u.a.r. and join the corresponding team
    reject all other requests
  end if
end if

```

number of rounds that is polynomial in the network size with high probability. The assumption that a stable matching exists is for ease of presentation, and all our results also hold for reaching approximate best matchings, by replacing $d(M)$ with $d(M) - d(M^*)$, where M^* is a best matching of G .

Given a network G , for every $t \geq 0$, let $M(t)$ be the matching of G at the beginning of round t , with deficit $d(M(t))$. The next property follows from the fact that leaders do not voluntarily disengage from the followers in their teams (and therefore the deficit of a leader increases of a unit only if the deficit of another leader decreases by one unit).

Property 1: For $t \geq 0$, $d(M(t))$ is non-increasing in t .

The next property follows from the assumption $c_\ell \geq 1, \forall \ell$.

Property 2: If G admits a stable matching, then $d(M(t)) \leq m$ for every $t \geq 0$.

We are now ready to state our main result. The maximum degree of the leaders is given by $\Delta = \max_{\ell \in L} |N_\ell|$.

Theorem 1: Let G be a network with m followers and which admits a stable matching. Let Δ be the maximum degree of the leaders. Fix $0 < \varepsilon < 1$, and let $c(\varepsilon) = \lfloor 1/\varepsilon \rfloor (pq)^{-\lfloor 1/\varepsilon \rfloor}$ and $\alpha(\varepsilon) = \lfloor 1/\varepsilon \rfloor$. Then, a $(1 - \varepsilon)$ -approximate stable matching of G is reached within $c(\varepsilon)\Delta^{\alpha(\varepsilon)}m^2$ rounds of the algorithm with probability at least $1 - me^{-m/8}$.

Example 1: If Δ is constant in the network size, then one can choose $\varepsilon = 1/\log m$, and Theorem 1 implies that a $(1 - 1/\log m)$ -approximate stable matching is reached in at most




Fig. 2. A deficit-decreasing path of length 5 is represented at the top of the figure: ℓ_0 is a poor leader, f_3 is an unmatched follower, and matching edges are highlighted. The path is “solved” by turning each matched edge into an unmatched edge and vice versa, as show at the bottom of the figure: ℓ_0 obtains an additional follower (and therefore its deficit decreases by a unit) and both ℓ_1 and ℓ_2 do not change their numbers of followers.

$m^{2+\log(\Delta/pq)} \log m$ rounds with probability that goes to one as $m \rightarrow \infty$.

To prove Theorem 1, we introduce the notion of *deficit-decreasing* path, that, in our setup, plays the role of the augmenting path in the context of one-to-one matching [9]. As we consider bipartite networks, a path alternates leaders and followers.

Definition 8 (Deficit-decreasing path): Given a matching M of G , a cycle-free path $P = \ell_0, f_1, \ell_1, \dots, f_k$ (of odd length $2k-1$) is a deficit-decreasing path relative to M if $(\ell_i, f_i) \in M$ for all $1 \leq i \leq k-1$, ℓ_0 is a poor leader, and f_k is an unmatched follower.

In words, a deficit-decreasing path starts at a poor leader with an edge not in M , ends at a follower that is not matched, and alternates edges in M and edges not in M . To justify the nomenclature, observe that, if $d(M) > 0$ and P is a deficit-decreasing path relative to M , a new matching M' such that $d(M') = d(M) - 1$ can be obtained by turning each unmatched edge of P into a matched edge, and vice versa. This is depicted in Fig. 2.

The proof of Theorem 1 builds on a technical result that, given a matching M with $d(M) \geq \varepsilon m$, guarantees the existence of a deficit-decreasing path of length at most $2\lceil 1/\varepsilon \rceil$. The existence of such a path allows us to bound the number of rounds needed for a one-unit reduction of the deficit. Our technical lemma extends a known result by Hopcroft and Karp [15, Theorem 1] given in the context of one-to-one matchings, but our proof is more subtle because, according to our setup, leaders can be matched to multiple followers and can have different size constraints c_ℓ . The symmetric difference of two sets A and B is defined as $A \oplus B = (A \setminus B) \cup (B \setminus A)$. Two paths are *follower-disjoint* if they do not share any follower (even though they might share some

leader).

Lemma 1: Let G admit a stable matching N . Let M be a matching of G with deficit $d(M) > 0$. Then, in $M \oplus N$ there are at least $d(M)$ follower-disjoint deficit-decreasing paths relative to M .

Proof: See Appendix A. ■

We make use of Lemma 1 through the following corollary.

Corollary 1: Let G be a network with m followers, admitting a stable matching N . Let M be a matching of G with deficit $d(M) \geq \varepsilon m$, for some $\varepsilon > 0$. Then, in $M \oplus N$ there exists a deficit-decreasing path relative to M of length at most $2\lfloor 1/\varepsilon \rfloor$.

Proof: By Lemma 1, if $d(M) \geq \varepsilon m$ and N is a stable matching of G , then in $M \oplus N$ we can choose εm follower-disjoint deficit-decreasing paths relative to M , whose cumulative length is at most $2m$ (as they do not share followers, and G is bipartite). Necessarily, one of them has length at most $2\lfloor 1/\varepsilon \rfloor$. ■

We are now ready to present the proof of Theorem 1.

A. Proof of Theorem 1

Let G be a network with n leaders and m followers, admitting a stable matching. Fix $0 < \varepsilon < 1$. Recall that we denote the matching at the beginning of round $t \geq 0$ by $M(t)$. For every $0 < x \leq 1$, let

$$\tau(x) = \min \{t \geq 0 : d(M(t)) < xm\}$$

be the first round at whose beginning the deficit is strictly smaller than xm . By Property 1, $\tau(x_2) \geq \tau(x_1)$ if $x_2 < x_1$. We are interested in bounding $\tau(\varepsilon)$.

Consider any round $t_1 \geq 0$. By Property 2, $d(M(t_1)) \leq m$, and therefore there exists $0 < \varepsilon' \leq 1$ such that $d(M(t_1)) = \varepsilon' m$ (we assume $\varepsilon' > 0$, as the case of $\varepsilon' = 0$ is trivial). Observe that $\tau(\varepsilon')$ can be equivalently defined as

$$\tau(\varepsilon') = \min \{t > t_1 : d(M(t)) < d(M(t_1))\},$$

that is, the first round after t_1 in which the deficit decreases. The following lemma bounds $\tau(\varepsilon')$.

Lemma 2: Let Δ be the maximum degree of the leaders in G . Let $d(M(t_1)) = \varepsilon' m$, for some $0 < \varepsilon' \leq 1$. Let $c(\varepsilon') = \lfloor 1/\varepsilon' \rfloor (pq)^{-\lfloor 1/\varepsilon' \rfloor}$ and $\alpha(\varepsilon') = \lfloor 1/\varepsilon' \rfloor$. Then

$$\Pr(\tau(\varepsilon') - t_1 \leq c(\varepsilon') \Delta^{\alpha(\varepsilon')} m) \geq 1 - e^{-m/8}.$$

Proof: Observe that $1 \leq \Delta \leq m$. For $t \geq 0$, let $2h(t) - 1 \geq 1$ be the length of the shortest deficit-decreasing path relative to $M(t)$. By Corollary 2, $2h(t) - 1 \leq 2\lfloor 1/\varepsilon' \rfloor$ for all $t_1 \leq t < \tau(\varepsilon')$. We distinguish the cases of $h(t_1) = 1$ and $h(t_1) > 1$.

Let $h(t_1) = 1$, that is, there is a deficit-decreasing path of length 1. We claim that with probability at least pq/Δ the deficit decreases by at least one unit during the next round of the algorithm. Consider a poor leader $\ell \in L$ such that ℓ, f_1 is a deficit-decreasing path of length 1 relative to $M(t_1)$. With probability at least p/Δ , ℓ attempts to recruit f_1 and, conditional on this event, f_1 considers ℓ 's proposal with probability q . It follows that the deficit decreases by one unit with probability at least pq/Δ . Observe that the number of poor leaders $\ell' \neq \ell$ such that $\ell' f_1$ is a deficit-decreasing path of length 1 does not influence the bound pq/Δ . Indeed, if multiple poor leaders propose to f_1 , then f_1 considers each proposal independently with probability q , and then chooses one uniformly at random.

Let now $h(t_1) > 1$, and let P be a deficit-decreasing path of length $2h(t_1) - 1 \geq 3$ ending at an unmatched follower f' . We look at how $h(t)$ evolves for $t_1 \leq t < \tau(\varepsilon')$, that is, until a deficit-decreasing path is “solved”. Fix any t such that $t_1 \leq t < \tau(\varepsilon')$. By the same argument as above, $h(t+1) = h(t) - 1$ with probability at least pq/Δ . Observe that, as long as $h(t) > 1$, f' remains unmatched during round t because any other deficit-decreasing path ending at f' (at the beginning of round t) has length at least $2h(t_1) - 1 \geq 3$.

Let $\gamma = pq/\Delta$ and $\alpha(\varepsilon') = \lfloor 1/\varepsilon' \rfloor$. By independence of successive rounds of the algorithm and the bound $h(t) \leq \alpha(\varepsilon')$ for all $t_1 \leq t < \tau(\varepsilon')$, with probability at least $\gamma^{\alpha(\varepsilon')}$, a sequence of $\alpha(\varepsilon') - 1$ rounds reduces the length of P (a shortest deficit-decreasing path ending at f') to 1 and then P gets “solved”. Consider phases of $\alpha(\varepsilon')$ consecutive rounds, and for phases $i = 0, 1, \dots$, let X_i be *iid* Bernoulli random variables with success probability $\gamma^{\alpha(\varepsilon')}$. We have that P gets solved (and the deficit decreases by one unit) in at most T consecutive phases) if and only if $\sum_{i=1}^T X_i \geq 1$. Letting $T = m\gamma^{-\alpha(\varepsilon')}$, a Chernoff bound implies that, for any $1/2 < \delta < 1$,

$$\begin{aligned} \Pr\left(\tau(\varepsilon') - t_1 > m\alpha(\varepsilon')\gamma^{-\alpha(\varepsilon')}\right) &= \Pr\left(\sum_{i=1}^T X_i = 0\right) \\ &< \Pr\left(\sum_{i=1}^T X_i < (1 - \delta)m\right) < e^{-m\delta^2/2} < e^{-m/8}. \end{aligned}$$

The lemma follows. ■

We use the result of Lemma 2 to prove the claim of the theorem. If the matching $M(0)$ at time

$t = 0$ is not stable then there exists $0 < \varepsilon_0 \leq 1$ such that $d(M(0)) = \varepsilon_0 m$, and therefore the number of rounds $\tau(\varepsilon_0)$ before a one-unit reduction of the deficit is at most $c(\varepsilon_0)\Delta^{\alpha(\varepsilon_0)}m$ with probability at least $1 - e^{-m/8}$. Similarly, if $M(\tau(\varepsilon_0))$ is not stable, then $d(M(\varepsilon_0)) = \varepsilon_1 m$ for some $\varepsilon_1 < \varepsilon_0$. Iterating the same argument, the deficit decreases by one unit in at most $c(\varepsilon_1)\Delta^{\alpha(\varepsilon_1)}m$ rounds with probability at least $1 - e^{-m/8}$, reaching a matching $M(\tau(\varepsilon_1))$ with $d(M(\varepsilon_1)) = \varepsilon_2 m < \varepsilon_1 m$. Given the target ε , as $\varepsilon_0 \leq 1$ and the deficit decreases by integer amounts, it follows that there exist $k \leq \varepsilon_0 m \leq m$ and a sequence

$$1 \geq \varepsilon_0 > \varepsilon_1 > \cdots > \varepsilon_{k-1} > \varepsilon_k > 0$$

such that $\varepsilon_{k-1} \geq \varepsilon > \varepsilon_k$. It follows that

$$\begin{aligned} \tau(\varepsilon) &\leq \sum_{i=0}^{k-1} c(\varepsilon_i)\Delta^{\alpha(\varepsilon_i)}m < c(\varepsilon_{k-1})\Delta^{\alpha(\varepsilon_{k-1})}m^2 \\ &< c(\varepsilon)\Delta^{\alpha(\varepsilon)}m^2 \end{aligned}$$

with probability at least $1 - me^{-m/8}$, by a union bound.

V. EXPONENTIAL CONVERGENCE

Theorem 1 gives a polynomial bound for reaching a $(1 - \varepsilon)$ -approximate stable matching for any constant $0 < \varepsilon < 1$ and any network. However, a similar guarantee cannot be derived for the case of a stable matching, as shown in this section through a counterexample. In particular, we define a sequence of networks of increasing size and maximum degree that diverges with the network size, and show that the number of rounds required to converge from an approximate matching to the stable matching is exponentially large in the network's size with high probability from an overwhelming fraction of the approximate matchings M such that $d(M) = 1$.

For $n \geq 1$, let $G_n = (L_n \cup F_n, E_n)$ be the network with n leaders and n followers (i.e., $L_n = \{\ell_1, \dots, \ell_n\}$ and $F_n = \{f_1, \dots, f_n\}$), with edges $E_n = \{(\ell_i, f_j) : 1 \leq i \leq n, j \leq i\}$, and team size constraints $c_\ell = 1$ for all $\ell \in L_n$, see Figure 3. G_n has maximum degree $\Delta = n$ and a unique stable matching given by $M_n^* = \{(\ell_i, f_i) : 1 \leq i \leq n\}$.

Theorem 2: For any matching M of G_n , let $\tau(M)$ denote the number of rounds to converge to the perfect matching when starting from M . Then, for any fixed constant $0 < \gamma < 1$, $\tau(M)$ is exponentially large in γn with high probability for a $1 - 2^{-(1-\gamma)n}$ fraction of all the matchings M such that $d(M) = 1$.




Fig. 3. The network G_n for $n = 6$. The matching M'_n is highlighted.

Here we only provide a sketch of the proof, whose details are presented in Appendix B. To get an understanding of the algorithm's dynamics, consider the matching

$$M'_n = \{(\ell_i, f_{i-1}) : 2 \leq i \leq n\},$$

highlighted in Figure 3 for the case of $n = 6$. Observe that $d(M'_n) = 1$ and, under M'_n , ℓ_1 is poor, and the remaining leaders are stable. According to the algorithm, ℓ_1 attempts to recruit f_1 (currently in ℓ_2 's team). If f_1 accepts, then ℓ_1 becomes stable and ℓ_2 becomes poor (and can in turn attempt to recruit either f_1 or f_2). After each round, there exists a unique poor leader until the stable matching is reached. The stable matching is reached when ℓ_{n-1} (ℓ_5 in Figure 3) becomes poor and then successfully recruits f_{n-1} (f_5 in Figure 3), and finally ℓ_n successfully recruits f_n .

In general, fix any matching M of G_n such that $d(M) = 1$. In M , there is a single poor leader ℓ_{i_0} and a single unmatched follower f_{i_K} . M is associated to a unique deficit-decreasing path $\ell_{i_0}, f_{i_0}, \dots, \ell_{i_{K-1}}, f_{i_{K-1}}, \ell_{i_K}, f_{i_K}$. If $K \geq 1$, we define $h(M) = i_{K-1}$ as the *height* of M .

Starting from M , for every $t < \tau(M)$, the matching $M(t)$ at round t has deficit $d(M(t)) = 1$ (by Property 1), a single poor leader $\ell_{i(t)}$, the single unmatched follower f_{i_K} and height $h(M(t)) = h(M) = i_{K-1}$. The stochastic process tracking the position of the poor leader $\ell_{i(t)}$ is not a classical random walk on $\{\ell_1, \dots, \ell_{i_K}\}$ and its transition probabilities at each round depend on the current matching. The time to reach stability is upper bounded by $\min\{t : i(t) = h(M)\}$, that is, the first round in which $\ell_{h(M)}$ becomes poor (since $\ell_{h(M)}$ can then match with $f_{h(M)}$ leaving ℓ_{i_K} poor, who can in turn match with the unmatched follower f_{i_K} , thus reaching the stable matching).

We can then consider a one-to-one correspondence between the matchings $M(t)$ reachable from M in which $i(t) \leq h(M)$ (note that $d(M(t)) = 1$ for each of them) and the nodes of a tree whose size is exponentially large in the height $h(M)$. In particular, we can show that the process $\{M(t) : t \geq 0, M(0) = M\}$ is equivalent to a classical random walk on the nodes of the tree, and that reaching the matching with $i(t) = h(M)$ corresponds to reaching the root of the tree. A random walk starting at any node of the tree visits the root after a number of steps that is exponentially large in the height $h(M)$ with high probability. Finally, the proof of Theorem 2 is completed by arguing that, for any constant $0 < \gamma < 1$, a $1 - 2^{-(1-\gamma)n}$ fraction of all matchings M of G_n such that $d(M) = 1$ have height $h(M) \geq \gamma n$.

VI. SIMULATIONS

In this section, the performance of our algorithm is further evaluated through simulation. In Figure 4, the algorithm's average convergence time on the sequence of networks G_n defined in Section V is shown (in logarithmic scale). On the one hand, the thick solid line suggests that the average number of rounds to reach a 0.9-approximate stable matching is upper bounded by a polynomial of small degree, consistently with Theorem 1. On the other hand, convergence to the stable matching requires an average number of rounds that grows exponentially in n (thin solid line), as predicted by Theorem 2. Moreover, the dotted line represents the average time after which all followers become matched, that grows slowly with n .

Figure 5 shows the algorithm's performance in reaching successively finer approximations of the best matching on random networks $G(n, m, \rho)$. Here, $G(n, m, \rho)$ refers to a random bipartite network with n leaders and m followers, in which each edge exists independently of the others with probability ρ (we fixed $\rho = 0.04$), and with constraint $c_\ell = \min\{m/n, |N_\ell|\}$ for each leader ℓ . For each of the (n, m) pairs that we considered, 20 random $G(n, m, \rho)$ were generated, and the algorithm was run 20 times on each. We observe that, consistently with Theorem 1, $\tau(\varepsilon)$ increases both when ε decreases (i.e., when a finer approximation is desired) and when the number m of followers increases. The plot visually suggests that a good solution is reached quickly, while most of the time is spent in the attempt of improving it to the best solution.




Fig. 4. Algorithm's convergence time on the sequence of networks G_n .




Fig. 5. Algorithm's average time to reach a $(1 - \varepsilon)$ -approximate best matching on random bipartite networks $G(n, m, \rho)$, for $\rho = 0.04$.

VII. DISCUSSION

The distributed algorithm we proposed, in which leaders and followers act according to simple local rules, is computationally tractable and allows us to derive performance guarantees in the form of theorems. Despite its simplicity, the algorithm is shown to reach an arbitrarily close approximation of a stable matching (or of a best matching) in polynomial time in any network. However, in general there can be an exponential gap between reaching an approximate solution

and a stable solution.

In the proposed algorithm, leaders do not communicate between each other, and only act in response to their own status and the status of their neighborhoods. The only collaboration between them consists in the fact that the leaders whose size constraints are satisfied do not attempt to recruit additional matched followers, and this is justified since recruiting more followers might be costly. How communication between leaders affects performance is an open question, as well as determining what amounts of communication and complexity are necessary to remove the exponential gap in the case of unbounded degree networks.

Finally, in Section V, we defined a sequence of networks in which the maximum degree of the leaders scales linearly with the network size. It would be interesting to understand whether a counterexample in which the maximum degree scales more slowly (e.g., logarithmically in the network size) could be derived.

APPENDIX

A. PROOF OF LEMMA 1

To prove the lemma we proceed as follows. First, we show that for each leader ℓ with deficit $d_\ell(M) > 0$ there are at least $d_\ell(M)$ deficit-decreasing paths that start at ℓ . Then, we argue that $d(M)$ follower-disjoint deficit-decreasing paths relative to M can be chosen.

Consider a leader ℓ that is poor under M , that is, $d_\ell(M) > 0$, and refer to Fig. 6. As ℓ has a team size constraint $c_\ell > 0$, there are exactly $c_\ell - d_\ell(M)$ followers that are matched with ℓ . As G admits a stable matching, there are $k \geq d_\ell(M)$ followers in N_ℓ that are not currently matched to ℓ (observe that $|N_\ell| \geq c_\ell$, where N_ℓ is the neighborhood of ℓ). Assume that, $h < d_\ell(M)$ of these followers are on distinct (follower-disjoint) deficit-decreasing paths starting at ℓ , denoted by P_1, \dots, P_h . Assume that the remaining $k - h > 0$ followers do not lay on follower-disjoint deficit-decreasing paths starting at ℓ , and denote them by f_1, \dots, f_{k-h} .

Consider a network G' obtained from G by removing all nodes on the paths P_1, \dots, P_h except from ℓ . Assume that ℓ has constraint $c'_\ell = c_\ell - h$ in G' , and consider the matching M' of G' given by

$$M' = M \setminus \{e : e \in (\cup_{i=1}^h P_i) \cap M\}.$$

Observe that $d_\ell(M)$ in G is equal to $d_\ell(M')$ in G' .




Fig. 6. A leader ℓ with constraint c_ℓ , degree $|N_\ell| \geq c_\ell$ and deficit $d_\ell(M)$. Matched edges are highlighted. ℓ is matched to exactly $c_\ell - d_\ell(M)$ followers (represented on the right). To prove Lemma 1 we assume that: all the remaining $k \geq d_\ell(M)$ followers in N_ℓ are matched in M , $h < k$ of them lay on deficit-decreasing paths starting at ℓ (denoted by P_1, \dots, P_h), and the other $k - h$ (denoted by f_1, \dots, f_k) do not lay on deficit-decreasing paths starting at ℓ .

In order to become stable, ℓ needs to match with at least one additional follower among $\{f_1, \dots, f_{k-h}\}$. We show that if ℓ matches with one of these followers then the deficit of another leader would eventually increase and G does not admit a stable matching, generating a contradiction.

Consider any follower $f' \in \{f_1, \dots, f_{k-h}\}$. Let ℓ' be the leader such that $(\ell', f') \in M$ (note that also $(\ell', f') \in M'$), and observe that if ℓ matches to f' then the deficit of ℓ' would increase to $d_{\ell'}(M) + 1$. The only way in which ℓ' can loose this additional unit of deficit (without increasing again the deficit of ℓ) is to match with a follower f'' such that $(\ell'', f'') \in M$ for some $\ell'' \neq \ell$, assuming that such f'' exists. Note that f'' is matched in M , otherwise ℓ, f', ℓ', f'' would be a deficit-decreasing path relative to M starting at ℓ , generating a contradiction (see Fig. 7). Therefore, the deficit of ℓ'' would increase to $d_{\ell''}(M) + 1 > 0$.

Observe that:

- (i) if $\ell'' = \ell$ then $\ell, f', \ell', f'', \ell$ would be a cycle and one of the leaders in the cycle must have deficit strictly greater than zero;
- (ii) if $\ell'' \neq \ell$ and ℓ'' is adjacent to no follower other than f'' then one of the leaders on the path $\ell, f', \ell', f'', \ell''$ must have deficit strictly greater than zero (see bottom of Fig. 7).

If instead ℓ'' is adjacent to a follower $f''' \neq f''$, then f''' is matched in M , as otherwise $\ell, f', \ell', f'', \ell'', f'''$ would be a deficit-decreasing path relative to M . Again, if either $\ell''' = \ell$, or




Fig. 7. If f'' is not matched then ℓ, f', ℓ', f'' would be a deficit-decreasing path (shown at the top of the figure, in which matching edges are highlighted), contradicting the assumption that P_1, \dots, P_h are the only deficit-decreasing paths relative to M starting at ℓ (or similarly, contradicting the assumption that in G' there are no deficit-decreasing paths relative to M'). Therefore, f'' is matched to a leader ℓ'' (the bottom of the figure represents the case of $\ell \neq \ell''$).

$\ell = \ell'$, or ℓ''' is not adjacent to an additional follower $f''' \neq f', f''$ (similarly to the cases (i) and (ii) above) then there would be a path or a cycle in which a leader must have deficit strictly greater than zero. By iterating this argument and considering all followers $f' \in \{f_1, \dots, f_{k-h}\}$, we have that ℓ cannot become stable without eventually increasing the deficit of at least another leader, contradicting the assumption that G admits a stable matching.

We showed that, letting N be a suitable stable matching of G , if a leader ℓ has deficit $d_\ell(M)$ then there are at least $d_\ell(M)$ deficit-decreasing paths in $M \oplus N$ relative to M that start at ℓ . To complete the proof of the lemma, we need to show that, among all the deficit-decreasing paths relative to M , we can choose $d(M)$ of them that are follower-disjoint.

Let \mathcal{P} denote any set of at least $d(M)$ deficit-decreasing paths relative to M such that at least $d_\ell(M)$ of them start at each ℓ (the argument above guarantees the existence of such a set \mathcal{P}). Assume by contradiction that we cannot choose \mathcal{P} such that all its elements are follower-disjoint. In particular, given \mathcal{P} , assume that there is a leader ℓ with deficit $d_\ell(M) > 0$ such that at most $d_\ell(M) - 1$ paths in \mathcal{P} starting at ℓ are follower-disjoint from all other paths in \mathcal{P} . In \mathcal{P} , there must be a deficit-decreasing path P starting at ℓ and a deficit-decreasing path P' starting at $\ell' \neq \ell$ that are not follower-disjoint (see Figures 8 and 9). It follows that if P is solved (by “flipping” matched edges into unmatched edges, and vice versa) then P' cannot be solved, and if P' is solved then P cannot be solved. Therefore, solving as many deficit-decreasing paths in \mathcal{P} as possible does not allow to obtain a stable matching from M . As the argument above holds for any set \mathcal{P} , we have that if G admits a stable matching then we can choose $d(M)$ follower-disjoint




Fig. 8. Given the matching at the top of the figure (matching edges are highlighted), assume that both ℓ_0 and ℓ_1 are poor, and that f_3 is unmatched. The deficit-decreasing paths $P = \ell_0, f_1, \ell_1, f_2, \ell_2, f_3$ and $P' = \ell_0, f_1, \ell_1, f_2$ intersect. If P' is solved (shown at the bottom of the figure), then P cannot be solved.




Fig. 9. If, under the matching highlighted in the figure, both ℓ and ℓ' are poor and f_1 is unmatched then there are two intersecting deficit-decreasing paths (one starting at ℓ and ending at f_1 , the other starting at ℓ' and ending at f_1).

deficit-decreasing paths relative to M .

B. PROOF OF THE EXPONENTIAL CONVERGENCE

Let \mathcal{M}_n be the set of all the matchings of G_n such that $d(M) = 1$. We proceed as follows. First, we show that each $M \in \mathcal{M}_n$ is uniquely identified by the set of the leaders that are not matched with “horizontal” edges (that defines a unique deficit-decreasing path), and we define the *height* $h(M)$ of M . Then we define trees T_m^* , $m \geq 1$ such that a random walk on T_m^* starting at any node different than the root hits the root after a number of steps that is exponentially large in m with high probability. We then argue that, when initialized with a matching $M \in \mathcal{M}_n$, the algorithm’s dynamics is equivalent to a random walk on the tree $T_{h(M)}^*$, and that reaching the stable matching of G_n corresponds to reaching the root of T_m^* (and therefore requires a number of rounds that is exponentially large in $h(M)$ with high probability). Finally, by a counting argument, we show that an overwhelming fraction of all the matchings in \mathcal{M}_n have height at least γn for any constant

$$0 < \gamma < 1.$$

A. The approximate stable matchings with $d(M) = 1$.

Matchings in \mathcal{M}_n enjoy some structural properties.

Lemma 3: Let $M \in \mathcal{M}_n$. The following properties hold.

- (1) There are a single poor leader $\ell_{i^*(M)}$ and a single unmatched follower $\ell_{j^*(M)}$ in M .
- (2) $1 \leq i^*(M) \leq j^*(M) \leq n$.
- (3) $(\ell_k, f_k) \in M$ for all $k < i^*(M)$ and all $k > j^*(M)$.
- (4) Let $\mathcal{J}(M) = \{j_1, j_2, \dots, j_K\}$ be the sorted set of indexes j such that $(\ell_j, f_j) \notin M$ for all $j \in \mathcal{J}(M)$. Then there exists a sorted set of indexes $\mathcal{I}'(M) = \{i_1, i_2, \dots, i_{K-1}\}$ such that $(\ell_{i_k}, f_{j_k}) \in M$ for all $k \in \{1, \dots, K-1\}$ and
 - (a) $j_1 = i^*(M)$.
 - (b) $j_K = j^*(M)$.
 - (c) $j_k < i_k$ for all $k \in \{1, \dots, K-1\}$.
 - (d) $i_k = j_{k+1}$ for all $k \in \{1, \dots, K-1\}$.

Proof: Property (1). As $d(M) = \sum_{\ell \in L} d_\ell(M) = 1$, there is a single poor leader $\ell_{i^*(M)}$ in M . As $c_\ell = 1$ for all $\ell \in L$, each leader $\ell \neq \ell_{i^*(M)}$ is matched to a single follower. It follows that there exists a unique unmatched follower $f_{j^*(M)}$.

Property (2). Suppose by contradiction that $i^*(M) > j^*(M)$. As $N_{\ell_{j^*(M)}} = \{f_1, \dots, f_{j^*(M)}\}$ and $f_{j^*(M)}$ is unmatched, leader $\ell_{j^*(M)}$ is matched to one among the followers in $\{f_1, \dots, f_{j^*(M)-1}\}$. As a consequence, the $j^*(M) - 1$ leaders in $\{\ell_1, \dots, \ell_{j^*(M)-1}\}$ are matched to at most $j^*(M) - 2$ out of the $j^*(M) - 1$ followers in $\{f_1, \dots, f_{j^*(M)-1}\}$, and one among them must be poor, contradicting Property (1). Therefore, $i^*(M) \leq j^*(M)$.

Property (3). We proceed by induction. If $i^*(M) > 1$, then $(\ell_1, f_1) \in M$ (since $N_{\ell_1} = \{f_1\}$ and ℓ_1 is matched with a follower). Assume that if $i^*(M) > j$ then $(\ell_k, f_k) \in M$ for all $k \leq j$. If $i^*(M) > j+1$, then, by the inductive assumption, ℓ_{j+1} can only be matched to f_{j+1} (as $N_{\ell_{j+1}} = \{f_1, \dots, f_{j+1}\}$). This shows that $(\ell_k, f_k) \in M$ for all $k < i^*(M)$. If $j^*(M) < n$ then $(\ell_n, f_n) \in M$ (as follower f_n is matched and ℓ_n is the only leader connected to f_n). Assume by induction that if $j^*(M) < j$ then $(\ell_k, f_k) \in M$ for all $k \geq j$. If $j^*(M) < j-1$, then, by the inductive assumption, f_{j-1} can only be matched to ℓ_{j-1} (since f_{j-1} is adjacent to $\ell_{j-1}, \dots, \ell_n$). This shows that $(\ell_k, f_k) \in M$ for all $k > j^*(M)$.

Property (4). If $K = 1$ then $M = \{(\ell_i, f_i) : i \neq i^*(M)\}$, $j^*(M) = i^*(M)$, and properties (4a-4d) trivially hold. Therefore, let $K \geq 2$. Let $\mathcal{J}(M) = \{j_1, j_2, \dots, j_K\}$ be the sorted set of indexes j such that follower f_j is not matched to leader ℓ_j in M (that is, $(\ell_j, f_j) \notin M$). Let $\mathcal{I}'(M) = \{i_1, i_2, \dots, i_{K'}\}$ be the sorted set of indexes i such that leader ℓ_i is matched to a follower different than f_i (that is, $(\ell_i, f_i) \notin M$ and ℓ_i is matched). By property (3), we have that $j_1 = i^*(M)$ and $j_K = j^*(M)$, therefore properties (4a) and (4b) follow. Hence, f_{j_1} is matched to ℓ_{i_1} for some $i_1 > j_1$. This implies that $(\ell_{i_1}, f_{i_1}) \notin M$ and that $i_1 \in \mathcal{J}(M)$. It is easy to see that $(\ell_k, f_k) \in M$ for all $k \in \{j_1 + 1, \dots, i_1 - 1\}$. This implies that there is no $j \in \mathcal{J}$ such that $j_1 < j < i_1$, and thus $i_1 = j_2$. Properties (3c) and (3d) follow by induction. Then, necessarily $K' = K - 1$. ■

Lemma 3 states that non-horizontal matching edges do not intersect. In particular, given a matching $M \in \mathcal{M}_n$, the set $\mathcal{I}'(M)$ represents the set of (the indexes of) the leaders that are matched with non-horizontal edges (see Figure 10). Let $i^*(M) = i_0$ (that is, the index of the unique unmatched leader). Lemma 3 implies that any matching $M \in \mathcal{M}_n$ is uniquely identified by the set $\mathcal{I}(M) = \{i_0\} \cup \mathcal{I}'(M)$. In particular, the following result holds.




Fig. 10. An example of a matching M of G_6 with $d(M) = 1$. M is uniquely determined by the set $\mathcal{I}(M) = \{2, 4, 6\}$, that encodes the following: ℓ_2 is not matched, ℓ_4 is matched with f_2 , ℓ_6 is matched with f_4 , f_6 is not matched. Also note that $P(M) = \ell_2, f_2, \ell_4, f_4, \ell_6, f_6$ is the unique deficit-decreasing path relative to M .

Lemma 4: There is a one-to-one correspondence between the matchings in \mathcal{M}_n and the elements of

$$\mathcal{S} = \{A : A \subseteq \{1, \dots, n\}\}.$$

Proof: Fix $K \leq n - 1$. Let $A = \{i_0, i_1, \dots, i_K\} \in \mathcal{S}$ such that $1 \leq i_0 < i_1 < \dots < i_K \leq n$. The matching $M \in \mathcal{M}_n$ such that $\mathcal{I}(M) = A$ is given by

$$M = \{(\ell_{i_{k+1}}, f_{i_k}) : 0 \leq k \leq K - 1\} \cup \{(\ell_k, f_k) : k \notin A\}.$$

Note that ℓ_{i_0} is poor and f_{i_K} is unmatched. It is easy to check that $M \in \mathcal{M}_n$. The converse is similar. ■

Recall that $M_n^* = \{(\ell_k, f_k) : 1 \leq k \leq n\}$ is the unique stable matching of G_n . Let $\mathcal{I}(M_n^*) = \emptyset$.

Lemma 3 and Lemma 4 imply that each matching $M \in \mathcal{M}$ is associated to a unique deficit-decreasing path.

Corollary 2: Fix $M \in \mathcal{M}_n$, and let $\mathcal{I}(M) = \{i_0, \dots, i_K\}$ as above. Then, the path

$$P(M) = \ell_{i_0}, f_{i_0}, \ell_{i_1}, f_{i_1}, \dots, \ell_{i_K}, f_{i_K}$$

is the unique deficit-decreasing path in $M \oplus M_n^*$ relative to M .

Proof: By Lemma 4, $\mathcal{I}(M)$ is uniquely determined. We first show that the set of edges defined by $P(M)$ is equal to $M \oplus M_n^*$. Observe that $M \setminus M_n^*$ is given by the non-horizontal edges in M , and that $M_n^* \setminus M$ is given by the horizontal edges that are not in M . Therefore, by Lemma 3,

$$M \setminus M_n^* = \{(\ell_{i_1}, f_{i_0}), (\ell_{i_2}, f_{i_1}), \dots, (\ell_{i_K}, f_{i_{K-1}})\},$$

$$M_n^* \setminus M = \{(\ell_{i_0}, f_{i_0}), (\ell_{i_1}, f_{i_1}), \dots, (\ell_{i_K}, f_{i_K})\},$$

and therefore the set of edges defined by $P(M)$ is equal to $M \oplus M_n^*$. To see that $P(M)$ is a deficit-decreasing path, observe that ℓ_{i_0} is poor, f_{i_K} is unmatched, all the other followers are matched, and $P(M)$ correctly alternates edges from $M_n^* \setminus M$ and from $M \setminus M_n^*$.

The uniqueness of $P(M)$ follows since there is no other way to connect the poor leader and the unmatched follower with a path. ■

Corollary 2 suggests that, given a matching $M \in \mathcal{M}_n$, the unique deficit-decreasing path $P(M)$ must be “solved” in order to reach the stable matching of G_n .

B. The tree T_m^*

Definition 9: Let T_1 be a labelled rooted tree with a singleton node with label 1. Inductively, for $2 \leq i \leq n - 1$, let T_i be the labelled rooted tree whose root is labelled with i and its children are T_1, \dots, T_{i-1} . We define T_m^* to be the tree with a root with label $m + 1$ whose only child is T_m (also see Figure 11). Let r^* denote the root of T_m^* .




Fig. 11. The three T_m^* for $m = 5$.

We show that the hitting time of r^* for a random walk on T_m^* starting at any node $u \neq r^*$ is exponential in m with high probability. For a node $u \neq r^*$, we call the edge that connects u to its parent u 's *exit* edge. For any subtree $T_i \subset T_m^*$, let Z_i be the random variable denoting the number of steps that it takes for a walk starting at the root of T_i to exit T_i (that is, to hit the parent of the root of T_i). The following lemma provides an exponential lower bound on Z_i

Lemma 5: There exist positive constants $\alpha, \gamma > 0$ such that, for all $i \geq 2$,

$$\Pr[Z_i \geq \gamma \cdot 2^{i/(\alpha \log^2 i)}] \geq 1 - \frac{1}{\log i}.$$

Proof: We proceed by induction on i . For convenience, define $g(i) = \alpha \log^2 i$ and $f(i) = \gamma \cdot 2^{i/g(i)}$ for some $\alpha, \gamma > 0$. For any $\alpha > 0$, we can choose $\gamma > 0$ such that $f(i) \leq 1$; therefore, as $Z_i \geq 1$ with probability 1, the claim holds trivially for any $i \leq i^*$, where i^* is a suitable large constant.

Now consider any $i \geq i^*$ and suppose the claim holds up to $i - 1$. Every time the walk is on the root of T_i , it exits T_i with probability $1/i$. Therefore, letting E_t be the event that the first t times the walk is on the root of T_i it does *not* exit T_i , we have $\Pr[E_t] \geq 1 - t/i$. Let $t = i/(2 \log i)$, and let D_j , $1 \leq j \leq t$, be the event that, when it is on the root of T_i for the j -th time, the walk moves to the root of one of the subtrees $T_{i-g(i)}, \dots, T_{i-1}$ and takes at least $f(i - g(i))$ steps to

exit that subtree. For $1 \leq j \leq t$, we have

$$\begin{aligned} \Pr[D_j \mid E_t] &\geq \frac{g(i)}{i} \cdot \Pr[Z_{i-g(i)} \geq f(i-g(i))] \\ &\geq \frac{g(i)}{i} \cdot \left(1 - \frac{1}{\log(i-g(i))}\right), \end{aligned}$$

by the induction hypothesis on $Z_{i-g(i)}$. Letting χ_j be the indicator function of the event D_j for $1 \leq j \leq t$, the probability that at least two of the events D_j happen, given E_t , is lower bounded by:

$$\begin{aligned} \Pr\left[\sum_{j=1}^t \chi_j \geq 2 \mid E_t\right] &\geq \Pr\left[\sum_{j=1}^{t/2} \chi_j \geq 1, \sum_{j=t/2+1}^t \chi_j \geq 1 \mid E_t\right] \\ &= \Pr\left[\sum_{j=1}^{t/2} \chi_j \geq 1 \mid E_t\right]^2. \end{aligned}$$

By union bound, we can write

$$\begin{aligned} &\Pr\left[\sum_{j=1}^{t/2} \chi_j \geq 1 \mid E_t\right] \\ &\geq 1 - \prod_{i=1}^{t/2} (1 - \Pr[D_j | E_t]) \\ &\geq 1 - \left(1 - \frac{g(i)}{i} \left(1 - \frac{1}{\log(i-g(i))}\right)\right)^{t/2} \\ &\geq 1 - \exp\left[-\frac{\alpha \log i}{4} \left(1 - \frac{1}{\log(i-g(i))}\right)\right] \geq 1 - \frac{1}{i^{\alpha/8}}, \end{aligned}$$

where the last step holds for i sufficiently large so that $\log(i-g(i)) \geq 2$. This implies that

$$\Pr\left[\sum_{j=1}^t \chi_j \geq 2 \mid E_t\right] \geq \left(1 - \frac{1}{i^{\alpha/8}}\right)^2 \geq 1 - \frac{2}{i^{\alpha/8}}.$$

Therefore, we conclude that

$$\begin{aligned} \Pr[Z_i \geq 2 \cdot f(i-g(i))] &\geq \Pr\left[\sum_{j=1}^t \chi_j \geq 2\right] \\ &\geq \Pr\left[\sum_{j=1}^t \chi_j \geq 2 \mid E_t\right] \Pr[E_t] \\ &\geq \left(1 - \frac{2}{i^{\alpha/8}}\right) \left(1 - \frac{t}{i}\right) \geq 1 - \frac{1}{\log i}, \end{aligned}$$

where the last step holds by choosing α sufficiently large. The claim follows since $2 \cdot f(i - g(i)) \geq f(i)$. \blacksquare

Note that a random walk starting at any node $u \neq r^*$ has to exit T_m before hitting r^* . Therefore, an application of Lemma 5 to T_m yields a lower bound to the hitting time of r^* when starting at any node $u \neq r^*$.

Corollary 3: The hitting time of r^* of a random walk starting at any node $u \neq r^*$ is $2^{\Omega(n/\log^2 n)}$ with high probability.

C. The dynamics of the algorithm starting from $M \in \mathcal{M}_n$

Fix $M \in \mathcal{M}_n$. By Lemma 4, each M is uniquely identified by the set $\mathcal{I}(M) = \{i_0, \dots, i_K\}$, for some $K \leq n - 1$. If $K = 0$ (that is, $|\mathcal{I}(M)| = 1$) then, according to the algorithm, the stable matching is reached in an average number $1/pq$ of steps, since the poor leader ℓ_{i_0} is adjacent to the unmatched follower f_{i_0} (note that there are exactly n matchings $M \in \mathcal{M}_n$ such that $|\mathcal{I}(M)| = 1$). In the remaining part of this section, we focus on the case of $K \geq 1$.

Consider a matching $M \in \mathcal{M}_n$ such that $K \geq 1$ (that is, $|\mathcal{I}(M)| > 1$). Define the *height* of M as $h(M) = i_{K-1}$, and observe that $h(M) \in \{1, \dots, n-1\}$. The stable matching M_n^* is eventually reached when the leader $\ell_{h(M)}$ becomes poor and recruits $f_{h(M)}$ (leaving ℓ_{i_K} poor), and then ℓ_{i_K} matches with f_{i_K} . Observe that, starting from the matching M , any matching M' such that $|\mathcal{I}(M')| > 1$ obtained before reaching the stable matching has height $h(M') = h(M)$. Note that the unique poor leader in M is ℓ_{i_0} , for $i_0 = \min \mathcal{I}(M)$. Also observe that $(\ell_{i_1}, f_{i_0}) \in M$, $i_1 = \min(\mathcal{I}(M) \setminus \min \mathcal{I}(M))$, and so on (note that f_{i_1} is the unmatched follower in M if and only if $P(M)$ has length 3).

We now analyze the dynamics of the algorithm when it is initialized with the matching M . For $t \geq 0$, let $M(t)$ be the matching at the beginning of round t . We have that $M(0) = M$. Let $\tau(M)$ be the number of rounds to reach the stable matching M_n^* . For ease of presentation, we set the probability parameters of the algorithms to¹ $p = q = 1$. For $0 \leq t \leq \tau(M)$, the matching $M(t)$ is uniquely identified by the set of sorted indexes $\mathcal{I}(M(t))$ or, equivalently, by the path $P(M(t))$. For ease of notation, let $\mathcal{I}(t) = \mathcal{I}(M(t))$, and $P(t) = P(M(t))$. We study the Markov process $\{\mathcal{I}(t), 0 \leq t \leq \tau(M)\}$. Observe that $\tau(M) = \min\{t : \mathcal{I}(t) = \emptyset\}$, as we defined $\mathcal{I}(M_n^*) = \emptyset$.

¹Setting $p=1$ means that, for $0 \leq t \leq \tau(M)$, the single poor leader in $M(t)$ always proposes to a follower. Setting $q = 1$ means that the follower always accepts the incoming request. Our result holds for any choice of p and q .

Fix $t \in \{0, \dots, \tau(M) - 1\}$. Letting $\mathcal{I}(t) = I \in \mathcal{S}$, the transition probabilities at time t are given by

$$\Pr(\mathcal{I}(t+1) = I' | \mathcal{I}(t) = I) = \frac{1}{\min I} \quad \text{if } I' \in \{I \cup \{k\} : k < \min I\} \cup \{I \setminus \{\min I\}\},$$

and 0 otherwise. To see this, observe that, at time t , letting $\mathcal{I}(t) = I$, the poor leader $\ell_{\min I}$ has degree $\min I$ and neighborhood $N_{\min I} = \{f_1, \dots, f_{\min I}\}$, and chooses one of the followers in $N_{\min I}$ uniformly at random. If $\ell_{\min I}$ chooses follower f_k for some $k < \min I$ then the leader ℓ_k becomes poor, since by property (3) of Lemma 3 ℓ_k was matched with f_k in $M(t)$, and we have that $\mathcal{I}(t+1) = I \cup \{k\}$, and the deficit-decreasing path associated to $M(t)$ is $P(t+1) = \ell_k, f_k, P(t)$. If instead $\ell_{\min I}$ chooses follower $f_{\min I}$ (matched with $\ell_{\min(I \setminus \min I)}$ in $M(t)$), then $\mathcal{I}(t+1) = I \setminus \{\min I\}$, and $P(t) = \ell_{\min I}, f_{\min I}, P(t+1)$ (and the stable matching is reached if $\mathcal{I}(t+1) = \emptyset$).

Let $\mathcal{R}(M)$ be the set of the matchings that can be reached from M (after one or multiple steps). It is easy to see that

$$\mathcal{R}(M) = \{M_n^*\} \cup \{M' \in \mathcal{M}_n : I(M') = A \cup \{h(M), i_K\}, A \subseteq \{1, \dots, h(M) - 1\}\},$$

where f_{i_K} is the unmatched follower in M . To be precise, $\mathcal{R}(M)$ should also include the matching M' such that $\mathcal{I}(M') = \{i_0\}$, but we identify M' and the stable matching M_n^* as (with $p = q = 1$) M_n^* is deterministically reached from M' in a single step. Observe that, for any $M' \in \mathcal{R}(M)$ such that $M' \neq M_n^*$, M' has height $h(M') = h(M)$.

Lemma 6: There is a one-to-one correspondence between the matchings in $\mathcal{R}(M)$ and the nodes of $T_{h(M)}$.

Proof: We define a mapping $\omega : \mathcal{R}(M) \rightarrow T_{h(M)}$ in the following way. Let $M_1 \in \mathcal{R}(M)$. Then M_1 is associated to a set $I \in \mathcal{S}$. We define $\omega(M_1)$ as the node of $T_{h(M)}$ with label $\min I$ and connected to the root with a path of nodes labelled by the sorted indices in $I \setminus \{\min I\}$. It is easy to verify that the defined mapping is a bijection. ■

Lemma 7: The stochastic process $\{\mathcal{I}(t) : 0 \leq t \leq \tau(M)\}$ is equivalent to a random walk on $T_{h(M)}$ starting at $\omega(M)$.

Proof: It suffices to show that the transition probabilities between two matchings $M_1, M_2 \in \mathcal{R}(M)$ are nonzero if and only if the nodes $\omega(M_1)$ and $\omega(M_2)$ are adjacent in $T_{h(M)}$. To prove the “only if” direction, assume that $M_1, M_2 \in \mathcal{R}(M)$ are such that there is a nonzero transition probability from M_1 to M_2 (and therefore from M_2 to M_1). Let $\mathcal{I}(M_1) = I_1$ and $\mathcal{I}(M_2) = I_2$. According to the transition probabilities given above, there are two possible cases. In the first

case, $I_2 = I_1 \cup \{k\}$ for some $k < \min I_1$, and $\omega(M_1)$ is a child of $\omega(M_1)$. In the second case $I_2 = I_1 \setminus \{\min I_1\}$ and $\omega(M_2)$ is the parent of $\omega(M_1)$. The proof of the other direction is similar. \blacksquare

As a consequence, reaching the stable matching of G_n starting from $M \in \mathcal{M}_n$ (for the algorithm) is equivalent to reaching the root of $T_{h(M)}$ starting from $\omega(M)$ (for a random walk on $T_{h(M)}$). By Corollary 3, $\tau(M)$ is exponentially large in $h(M)$ with high probability. To complete the proof of the Theorem, we show that, for any constant $0 < \gamma < 1$, a $1 - 2^{-(1-\gamma)n}$ fraction of the matchings $M \in \mathcal{M}_n$ have $h(M) \geq \gamma n$. This is done through a counting argument.

D. The fraction of the matchings $M \in \mathcal{M}_n$ such that $h(M) \geq \gamma n$

Let N be the number of the matchings in \mathcal{M}_n . Fixed a constant $0 < \gamma < 1$, let $\mathcal{M}_\gamma = \{M \in \mathcal{M}_n : h(M) < \gamma n\}$ and let N_γ be its cardinality. Recall that there are $N(0) = n$ matchings $M \in \mathcal{M}_n$ such that $|\mathcal{I}(M)| = 1$. For the case $|\mathcal{I}(M)| > 1$, M has height $h(M) \in \{1, \dots, n-1\}$. For $k = 1, \dots, n-1$, let $N(k)$ be the number of matchings $M \in \mathcal{M}_n$ such that $h(M) = k$. It follows that

$$N = \sum_{k=0}^{n-1} N(k), \quad N_\gamma \leq \sum_{k=0}^{\lceil \gamma n \rceil - 1} N(k).$$

Lemma 8: $N(k) = (n-k)2^{k-1}$ for all $k = 1, \dots, n-1$.

Proof: Fix $k \in \{1, \dots, n-1\}$. We count the number the matchings $M \in \mathcal{M}$ such that either (i) ℓ_k is the last *stable* leader on the path $P(M)$ (if $P(M)$ has length strictly greater than 3), or (ii) ℓ_k is the poor leader in M and $P(M)$ has length 3.

The unmatched follower under such a matching M is $\ell_{j^*(M)}$ for some $j^*(M) \in \{k+1, \dots, n\}$. By Lemma 3, all the leaders ℓ_i for $k < i \neq j^*(M)$ are matched with horizontal edges. Thus, there are $n-k$ possible choices for the leader ℓ_i matched with a non-horizontal edge that has the largest index i , and

$$N(k) = (n-k)D(k),$$

where $D(k)$ is the number of matchings $M' \in \mathcal{M}_{k+1}$ ² such that $h(M') = k$ (note that the leader matched with non-horizontal edge in M' with the largest index is ℓ_{k+1}). Such a matching M' of

²Similarly to \mathcal{M}_n , we define \mathcal{M}_{k+1} as the set of the matchings M' of G_{k+1} such that $d(M') = 1$.

G_{k+1} has a single poor leader ℓ_{i_0} for some $i_0 \in \{1, \dots, k\}$. Therefore we have that

$$D(k) = \sum_{i_0=1}^k F(i_0),$$

where $F(i_0)$ is the number of matchings $M' \in \mathcal{M}_{k+1}$ such that ℓ_{i_0} is poor and $h(M') = k$. If $i_0 = k$, then it is easy to see that $F(k) = 1$ (this is the case in which ℓ_k is poor and $P(M')$ has length 3). Consider now $1 \leq i_0 < k$. By Lemma 3, all leaders ℓ_i , $i < i_0$ are matched in M' with horizontal edges, ℓ_{i_0} is not matched in M' , and $(\ell_{k+1}, f_k) \in M'$. This implies that, to count $F(i_0)$, we have to count all possible sets $\mathcal{I} \subseteq \{i_0 + 1, \dots, k-1\}$, each corresponding to a matching M' in which ℓ_i is matched with a non-horizontal edge for all $i \in \mathcal{I}$ (note that ℓ_k and ℓ_{k+1} are already constrained to have a non-horizontal matching edge). We have that

$$F(i_0) = 2^{k-i_0-1},$$

for all $i_0 \in \{1, \dots, k-1\}$. It follows that

$$D(k) = \sum_{i_0=1}^k F(i_0) = 1 + \sum_{i_0=1}^{k-1} 2^{k-i_0-1} = 2^{k-1},$$

and therefore that $N(k) = (n-k)2^{k-1}$. ■

We now show that for any constant γ , the fraction of matchings $M \in \mathcal{M}_n$ such that $h(M) < \gamma n$ rapidly goes to zero as $n \rightarrow \infty$.

Lemma 9: Fix $0 < \gamma < 1$. Then, $N_\gamma/N = \Theta(2^{-(1-\gamma)n})$.

Proof: We first compute N .

$$N = \sum_{i=0}^{n-1} N(i) = n + \sum_{i=1}^{n-1} (n-i)D(i+1) = n + n \sum_{i=0}^{n-2} 2^i - \sum_{i=1}^{n-1} i2^{i-1}.$$

The second sum can be shown (e.g. by induction) to be equal to $(n-1) + (n-2)(2^{n+1} - 1)$. Therefore,

$$N = n + n(2^{n-1} - 1) - (n-1) - (n-2)(2^{n+1} - 1) = 2^n - (n+1) + n.$$

Similarly, we have that,

$$N_\gamma \leq \sum_{i=0}^{\lceil \gamma n \rceil - 1} N(i) = 2^{\lceil \gamma n \rceil} - (\lceil \gamma n \rceil + 1) + n.$$

Therefore, the fraction of matchings in \mathcal{M}_n with height $h(M) < \gamma n$ is $N_\gamma/N = \Theta(2^{-(1-\gamma)n})$. ■

REFERENCES

- [1] S. Abdallah and V. Lesser. Organization-based cooperative coalition formation. In *IEEE Conference on Intelligent Agent Technology*, pages 162–168, 2004.
- [2] I. Ashlagi, M. Braverman, and A. Hassidim. Matching with couples revisited. In *ACM EC*, pages 335–336, 2011.
- [3] S. Berman, A. Halasz, M.A. Hsieh, and V. Kumar. Optimized stochastic policies for task allocation in swarms of robots. *IEEE Trans. on Robotics*, 25(4):927–937, 2009.
- [4] V.D. Blondel, J.M. Hendrickx, A. Olshevsky, and J.N. Tsitsiklis. Convergence in multiagent coordination, consensus, and flocking. In *IEEE CDC*, pages 2996–3000, 2005.
- [5] F. Bullo, J. Cortés, and S. Martinez. *Distributed control of robotic networks: a mathematical approach to motion coordination algorithms*. Princeton Univ Pr, 2009.
- [6] A. Cenedese, F. Cerruti, M. Fabbro, C. Masiero, and L. Schenato. Decentralized task assignment in camera networks. In *IEEE CDC*, pages 126–131, 2010.
- [7] K.S. Cook, R.M. Emerson, M.R. Gillmore, and T. Yamagishi. The distribution of power in exchange networks: Theory and experimental results. *American Journal of Sociology*, pages 275–305, 1983.
- [8] L. Coviello, M. Franceschetti, M.D. McCubbins, R. Paturi, and A. Vattani. Human matching behavior in social networks: an algorithmic perspective. *PLoS One*, 2012.
- [9] R. Diestel. Graph theory. *Graduate texts in mathematics*, pages 24–26, 2005.
- [10] D. Enemark, M. McCubbins, R. Paturi, and N. Weller. Does more connectivity help groups to solve social problems? In *ACM EC*, 2011.
- [11] M. Franceschetti and J. Bruck. A group membership algorithm with a practical specification. *IEEE Trans. Parallel and Distributed Systems*, 12(11):1190–1200, 2001.
- [12] D. Gale and L.S. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.
- [13] M.E. Gaston and M. desJardins. Agent-organized networks for dynamic team formation. In *ACM AAMAS*, pages 230–237, 2005.
- [14] J.W. Hatfield and S.D. Kominers. Multilateral matching. In *ACM E-Commerce*, pages 337–338, 2011.
- [15] J.E. Hopcroft and R.M. Karp. A $n^{5/2}$ algorithm for maximum matchings in bipartite. In *IEEE SSAT*, pages 122–125, 1971.
- [16] B. Horling and V. Lesser. A survey of multi-agent organizational paradigms. *The Knowledge Engineering Rev.*, 19(04):281–316, 2004.
- [17] S. Judd, M. Kearns, and Y. Vorobeychik. Behavioral dynamics and influence in networked coloring and consensus. *PNAS*, 107(34):14978–14982, 2010.
- [18] Y. Kanoria, M. Bayati, C. Borgs, J. Chayes, and A. Montanari. Fast convergence of natural bargaining dynamics in exchange networks. In *ACM STOC*, pages 1518–1537, 2011.
- [19] M. Kearns, S. Judd, J. Tan, and J. Wortman. Behavioral experiments on biased voting in networks. *PNAS*, 106(5):1347, 2009.
- [20] M. Kearns, S. Suri, and N. Montfort. An experimental study of the coloring problem on human subject networks. *Science*, 313(5788):824, 2006.
- [21] J. Kleinberg. Cascading behavior in networks: Algorithmic and economic issues. *Algorithmic game theory*, pages 613–632, 2007.
- [22] J. Kleinberg and É. Tardos. Balanced outcomes in social exchange networks. In *ACM STOC*, pages 295–304, 2008.

- [23] Z. Lotker, B. Patt-Shamir, and S. Pettie. Improved distributed approximate matching. In *ACM SPAA*, pages 129–136, 2008.
- [24] N.A. Lynch. *Distributed algorithms*. Morgan Kaufmann, 1996.
- [25] M.D. McCubbins, R. Paturi, and N. Weller. Connected coordination: Network structure and group coordination. *American Politics Research*, 37:899, 2009.
- [26] A. Montanari and A. Saberi. The spread of innovations in social networks. *PNAS*, 107(47):20196, 2010.
- [27] B.J. Moore and K.M. Passino. Distributed task assignment for mobile agents. *IEEE Trans. on Aut. Control*, 52(4):749–753, 2007.
- [28] A. Nedić and A. Ozdaglar. Convergence rate for consensus with delays. *Journal of Global Optimization*, 47(3):437–456, 2010.
- [29] A. Nedić, A. Ozdaglar, and P.A. Parrilo. Constrained consensus and optimization in multi-agent networks. *IEEE Trans. on Aut. Control*, 55(4):922–938, 2010.
- [30] A. Olshevsky and J.N. Tsitsiklis. Convergence rates in distributed consensus and averaging. In *IEEE CDC*, pages 3387–3392, 2006.
- [31] F. Pasqualetti, S. Martini, and A. Bicchi. Steering a leader-follower team via linear consensus. *Hybrid Systems: Computation and Control*, pages 642–645, 2008.
- [32] S. Pettie and P. Sanders. A simpler linear time $2/3$ -approximation for maximum weight matching. *Information Processing Letters*, 91(6):271–276, 2004.
- [33] A.R. Rahmani, M. Ji, M. Mesbahi, and M.B. Egerstedt. Controllability of multi-agent systems from a graph-theoretic perspective. *SIAM Journal on Control and Optimization*, 48(1):162–186, 2010.
- [34] A.E. Roth. The evolution of the labor market for medical interns and residents: a case study in game theory. *The Journal of Political Economy*, pages 991–1016, 1984.
- [35] A.E. Roth and M.A.O. Sotomayor. *Two-sided matching: A study in game-theoretic modeling and analysis*. Cambridge Univ Pr, 1992.
- [36] A. Tahbaz-Salehi and A. Jadbabaie. A necessary and sufficient condition for consensus over random networks. *IEEE Trans. on Aut. Control*, 53(3):791–795, 2008.
- [37] A. Tahbaz-Salehi and A. Jadbabaie. Consensus over ergodic stationary graph processes. *IEEE Trans. on Aut. Control*, 55(1):225–230, 2010.
- [38] H.G. Tanner. On the controllability of nearest neighbor interconnections. In *IEEE CDC*, pages 2467–2472, 2004.
- [39] P. Tosic and G. Agha. Maximal clique based distributed group formation for autonomous agent coalitions. In *Coalitions and Teams Workshop, AAMAS*, 2004.
- [40] L. Vig and J.A. Adams. Multi-robot coalition formation. *IEEE Transactions on Robotics*, 22(4):637–649, 2006.
- [41] H.P. Young. The diffusion of innovations in social networks. *The Economy As an Evolving Complex System III*, page 267, 2006.